Estimating Credit Contagion in a Standard Factor Model

Daniel Rösch* and Birker Winterfeldt†
University of Regensburg

First Version: September 8, 2006
This Version: January 30, 2007

Abstract

State-of-the-art credit risk portfolio models and the New Basel Capital Accord consider only symmetric dependencies between borrowers in a portfolio, such as correlations. Recently, asymmetric dependencies have been introduced by Davis/Lo (2001) among others. However, statistical estimation techniques and empirical evidence on contagion is still rather scarce. The present paper provides a simple credit risk portfolio model extension to credit contagion and shows how its parameters can be easily estimated and tested. We apply our methodology to a dataset provided by Moody's Investor Services and find significant contagion effects. By stress testing we show how contagion can seriously affect credit losses.

Key Words: Credit Risk Models, Credit Contagion, Stress Testing

JEL Classification: G20, G28, C51

*Dr. Daniel Rösch (corresponding author), Department of Statistics, Faculty of Business and Economics, University of Regensburg, 93040 Regensburg, Germany
Phone: +49-941-943-2752, Fax : +49-941-943-4936
Email: daniel.roesch@wiwi.uni-regensburg.de, Internet: http://www.danielroesch.de

†Birker Winterfeldt, Department of Statistics, Faculty of Business and Economics, University of Regensburg, 93040 Regensburg, Germany

1 Introduction

Among the most important positions on the asset side of a financial institution's balance sheet are credit risky securities, and a major task for risk managers and analytics is the appropriate modeling and forecasting of the inherent credit risk. Typically, banks and other institutions apply credit risk models for this purpose, either purchased from a vendor such as CreditMetrics or CreditRisk+ or internally developed, see e.g. Finger (1998), Credit Suisse First Boston (1998), or Bluhm/Overbeck/Wagner (2003) for overviews. Credit risk models use borrower default probabilities, losses given default, exposure sizes, correlations among borrowers, and other parameters as input variables and derive forecasts for loss or market value distributions using analytical approaches or simulation techniques.

A common feature of most state-of-the-art credit risk models is the treatment of the dependencies between borrowers in a symmetric way, e.g. via correlations, dependencies to common risk factors or more general dependency measures in a copula framework. As many studies show, the strength of dependency crucially affects the shape of the derived distributions and key risk figures such as Value-at-Risk or Expected Shortfall, see Frey/McNeil/Nyfeler (2001) and Hamerle/Rösch (2005). This symmetric view can also be found in the New Basel Capital Accord where a key driver for banks' regulatory capital is correlation, see Basel Committee on Banking Supervision (2004).

More recently, researchers and practitioners argued that besides symmetric dependencies due to common risk factors asymmetric dependencies might exist as well. These effects are often called infection or contagion effects because the dependency works into one direction only. For example, the default of a large automobile company might cause financial difficulties of its suppliers, while conversely the automobile company might not be affected by a default of one of its suppliers.

One of the first models which explicitly model credit contagion is due to Davis/Lo (2001, DL hereafter). They consider a portfolio where the default of any company may infect any other company in the portfolio. Under this kind of contagion the portfolio loss distribution can easily be derived. An extension of the model is provided by Egloff/Leippold/Vanini (2004) who use neural-network-like connections between borrowers which allow for a variety of inter-firm infections. However, this model cannot be as easily applied in practice as the simple DL model because detailed information regarding the microstructural depencies is needed. Another model is due to Neu/Kühn (2004, NK hereafter) who incorporate contagion effects into a CreditMetrics-like credit risk model, thereby linking contagion with state-of-the-art models.

The DL and NK models build the starting point for our analysis. We first discuss some limitations of the former model when it is applied to real-world data. Then, we develop a simple contagion extension of factor models which are used in most credit risk models. The extension is similar to NK, however it is not as portfolio constrained as their approach.

Our main contribution consists in deriving a framework for empirical estimation and calibration of contagion effects. We apply our methodology to rating data from Moody's Investor Services and conduct some stress testing. Our main findings are that contagion effects are significant and can seriously affect loss distributions.

The rest of the paper is organized as follows. Section 2 provides a short discussion of the DL model. In section 3 we describe our contagion model extension of usual credit factor models and section 4 derives the estimation framework. Section 5 provides a description of the data used in our analysis. In section 6 we present our empirical results. Section 7 concludes.

2 Reviewing the Davis/Lo-Model

Davis/Lo (2001) were among the first to model contagion effects in a bond portfolio. They assume that any bond may default either directly or may be infected by any defaulting bond in the portfolio. p denotes the probability of a direct default, n the number of bonds in the portfolio, and q is the probability with which a defaulting bond infects another bond. Then the expected default rate $E[DR]$ in a portfolio equals

$$E[DR] = 1 - (1-p)(1-pq)^{n-1}. \tag{1}$$

As can be taken from formula (1) the expected default rate depends not only on the parameters p and q but also on the portfolio size n. Therefore, these parameters cannot be interpreted on a stand alone basis. The higher the number of firms in the portfolio is, the higher is the probability of any firm to be infected and thus (ceteris paribus) the expected default rate of the portfolio.

Despite its intuitiveness and simplicity, the model has some limitations when applied to real-world portfolios. First of all, it is assumed that all bonds within a portfolio may be infected by other bonds in the same portfolio only. In the real world contagion effects surely cross portfolio borders. Bonds outside the portfolio are not considered in the model. When two identical portfolios consisting of 50 bonds each (e.g. $p = 0.05$ and $q = 0.05$, expected default rate according to model: 15.97%) are combined into a single portfolio ($n = 100$), the model results in a jump of the expected default rate to 25.85% for constant parameters p and q. This result indicates that the model holds only in a portfolio of firms with no connections to firms outside the portfolio.

Another shortcoming becomes apparent when estimating the unknown parameters p and q of a portfolio. The parameters can easily be estimated from historical data e.g. by maximum-likelihood estimation. However, if the portfolio size varies over time the resulting parameter estimates cannot be interpreted. Another constraint is the mathematical operability. The distribution function requires the computation of a sum of binomial coefficients which becomes cumbersome for large portfolios.

In the following, we propose an alternative credit contagion model which can be estimated from historical data, is mathematically simple, and can be applied to portfolios of every size.

3 The Models

3.1 Standard Credit Risk Factor Model Specification

Our model is an extension of one of the most popular credit factor representations as it is used in CreditMetrics and also in the Basel II Capital Accord. We assume a default mode framework with a discrete-time horizon. Moreover, it is assumed that borrowers can be grouped into M segments, such that all borrowers within a segment are homogeneous with respect to parameters of their default processes, for instance rating grades. Consider a continuous variable $R_{i,t}$ of borrower i in time period t ($i \in N_t$, $t = 1, \ldots, T$) which may be interpreted as some creditworthiness index, e.g. the return on the firm's assets. Then the credit default event is modeled as the event that the creditworthiness of the firm crosses some threshold $c_{S(i)}$, i.e.

$$R_{i,t} < c_{S(i)} \quad \Leftrightarrow \quad D_{i,t} = 1, \tag{2}$$

where

$$D_{i,t} = \begin{cases} 1 & \text{borrower } i \text{ defaults in period } t \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

is the default indicator ($i \in N_t$, $t = 1, \ldots, T$). $S(i)$ denotes the segment borrower i belongs to ($S(\cdot) : \{1, \ldots, N_t\} \to \{1, \ldots, M\}$), e.g. a rating grade.

In the CreditMetrics (and the Basel II) framework the creditworthiness indexes are assumed to follow Gaussian copulas, that is

$$R_{i,t} = \omega_{S(i)} \cdot F_{S(i),t} + \sqrt{1 - \omega_{S(i)}^2} \cdot U_{i,t} \tag{4}$$

where $F_{S(i),t} \sim N(0,1)$ and $U_{i,t} \sim N(0,1)$ are both normalized i.i.d. random variables and independent from each other ($i \in N_t$, $t = 1, \ldots, T$). $F_{S(\cdot),t}$ is a systematic risk factor which drives all credit qualities within the segment jointly while $U_{i,t}$ are idiosyncratic, borrower-specific risk factors. $\omega_{S(\cdot)}$ is the factor loading of the systematic factor. The vector of systematic risk drivers $\boldsymbol{F}_t = (F_{1,t}, \ldots, F_{M,t})$ is normally distributed with mean

$$E(\boldsymbol{F}_t) = \mathbf{0} \tag{5}$$

and covariance matrix

$$\mathrm{Cov}(\boldsymbol{F}_t) = \begin{pmatrix} 1 & \sigma_{1,2} & \cdots & \sigma_{1,M} \\ \sigma_{1,2} & 1 & \cdots & \vdots \\ \vdots & \vdots & \ddots & \sigma_{M-1,M} \\ \sigma_{1,M} & \cdots & \sigma_{M-1,M} & 1 \end{pmatrix}, \tag{6}$$

$(t = 1, \ldots, T)$. $\sigma_{l,l'}$ is the correlation between the systematic factors of segments l and l'. The correlation between the creditworthiness indices of two firms is given by

$$\mathrm{Corr}\left(R_{i,t}, R_{j,t}\right) = \omega_{S(i)} \cdot \omega_{S(j)} \cdot \sigma_{S(i),S(j)}. \tag{7}$$

Given a realization of the systematic risk factor, the conditional probability of default is

$$\pi_i\left(f_{S(i),t}\right) = P\left(R_{i,t} < c_{S(i)} | F_{S(i),t} = f_{S(i),t}\right) = \Phi\left(\frac{c_{S(i)} - \omega_{S(i)} \cdot f_{S(i),t}}{\sqrt{1 - \omega_{S(i)}^2}}\right) \tag{8}$$

with expectation (the "probability of default")

$$\pi_i = \int_{-\infty}^{\infty} \Phi\left(\frac{c_{S(i)} - \omega_{S(i)} \cdot f_{S(i),t}}{\sqrt{1 - \omega_{S(i)}^2}}\right) d\Phi\left(f_{S(i),t}\right) = P\left(R_{i,t} < c_{S(i)}\right) = \Phi\left(c_{S(i)}\right). \tag{9}$$

$\Phi(\cdot)$ denotes the cumulative distribution function of the standard normal distribution.

3.2 Credit Contagion

We now assume that besides a set of firms which follows the above standard specification a distinct set of firms exists. These firms additionally depend upon the first group via credit contagion. That is, default events in the first group infect firms in the second group to an increase of their default probabilities but not vice versa. This contagion channel may be within the same segment as defined by the standard factor model, or we may also think of one-way infection which crosses these segmentations and follows a distinct hierarchy. For example, we may think of the standard credit factor model segmenting by rating grade, and credit contagion within an industry sector. We may therefore separate firms within a homogeneous rating grade additionally into different industry sectors and divide each sector into those firms which infect ("I-firms") other firms and those which are affected by contagion ("C-firms"). Figure 1 illustrates the contagion channels.

The process for the infecting firms is then given as above by

$$R_{i,t}^I = \omega_{S(i)} \cdot F_{S(i),t} + \sqrt{1 - \omega_{S(i)}^2} \cdot U_{i,t}. \tag{10}$$

Let $K(j)$ denote the contagion segment borrower j belongs to, e.g. an industry sector $(K(\cdot) : \{1, \ldots, N_t\} \to \{1, \ldots, K\})$. Let further $I_t\left[m, K(j)\right] = \{i | S(i) = m, K(i) = K(j), \text{infecting}\}$ be the set of infecting firms belonging both to rating grade m and industry segment $K(j)$ at time t, and $I_{K(j),t}$ be their number.

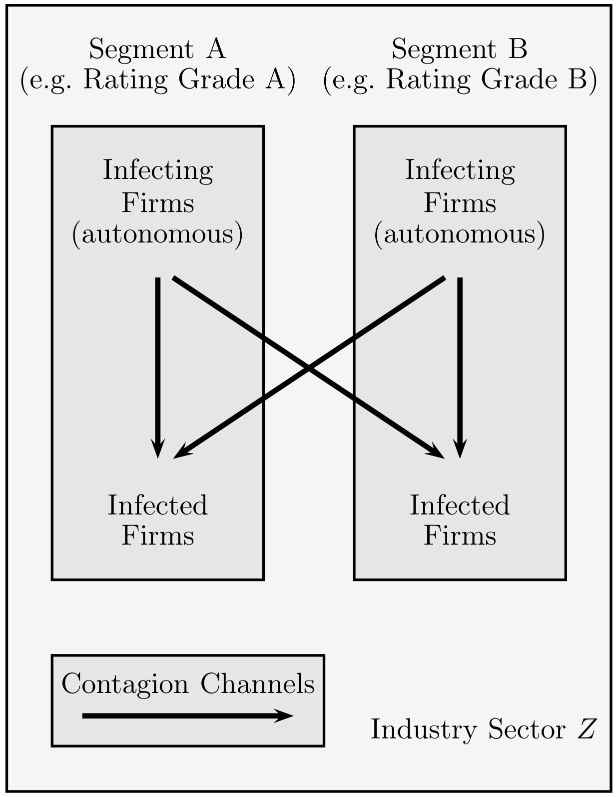

Figure 1: Contagion channels

Then we extend the standard factor specification for the infection of the second group of firms to

$$
\begin{aligned}
R^C_{j,t} &= \omega_{S(j)} \cdot F_{S(j),t} + \sqrt{1-\omega^2_{S(j)}} \cdot U_{j,t} + \beta \cdot \frac{\sum_{s\in I_t[m,K(j)]} D^I_{s,t}}{I_{K(j),t}} \\
&= \omega_{S(j)} \cdot F_{S(j),t} + \sqrt{1-\omega^2_{S(j)}} \cdot U_{j,t} + \beta \cdot \frac{D^I_{K(j),t}}{I_{K(j),t}}. \qquad (11)
\end{aligned}
$$

$D^I_{K(j),t} = \sum_{s\in I_t[m,K(j)]} D^I_{s,t}$ is the number of defaulting infecting firms in segment $K(j)$. β denotes an unknown coefficient which measures the impact of contagion on the default probability. The effect of contagion on firm j is then β times the default rate of the infecting firms within industry $K(j)$. Note that if β equals zero there is no contagion at all and the model reduces to the standard factor model.

While the probabilities of default for the infecting firms are still given by the "autonomous" probabilities (8) and (9), the probabilities of default of the infected firms now depend additionally on the default rate of the contaminating firms. Conditional on the risk factors and the number of defaulting infectors $D^I_{K(j),t} = d^I_{K(j),t}$ one obtains the conditional probability

$$\pi_j^C\left(f_{S(j),t}, d^I_{K(j),t}(\boldsymbol{f}_t)\right) = \Phi\left(\frac{c_{S(j)} - \omega_{S(j)} \cdot f_{S(j),t} - \beta \cdot \frac{d^I_{K(j),t}(\boldsymbol{f}_t)}{I_{K(j),t}}}{\sqrt{1-\omega^2_{S(j)}}}\right). \qquad (12)$$

To the extent that either β is zero or there are no defaulting contaminators, or both, the infected probability reduces to the "autonomous" conditional probability $\pi_i\left(f_{S(i),t}\right)$. The higher the absolute value of β and the number of defaults, the higher the infected conditional probability will be ceteris paribus.

4 Model Estimation

After outlining the model framework we want to estimate the models from observed data. We suggest a maximum-likelihood approach. For standard factor models, i.e. without contagion, this approach has been employed e.g. by Gordy/Heitfield (2000) for the Gaussian specification of Credit Metrics or Hamerle/Rösch (2007) for other Bernoulli mixture models, as used by CreditRisk+ or CreditPortfolioView.

The likelihood is based on the respective probabilities of observing particular numbers of defaulted infectors and defaulted infected firms. Let $I_t\left[m,k\right] = \{i|S(i)=m, K(i)=k, \text{infecting}\}$ be the set of infecting firms belonging both to rating grade m and industry segment k at time t and $I_{m,k,t}$ be their number.

$$D^I_{i,t} = \begin{cases} 1 & \text{borrower } i \text{ defaults in period } t \\ 0 & \text{otherwise} \end{cases} \qquad (13)$$

be the default indicator variable. If we condition on the common systematic factor $f_{m,t}$, the probability of observing $d^I_{m,k,t} = \sum_{i \in I_t[m,k]} d^I_{i,t}$ defaulting firms from $I_t\left[m,k\right]$ is given by

$$P^I_{m,k,t}(f_{m,t}) = P(d^I_{m,k,t}|f_{m,t}) = \binom{I_{m,k,t}}{d^I_{m,k,t}} \cdot \pi(f_{m,t})^{d^I_{m,k,t}} \cdot (1-\pi(f_{m,t}))^{(I_{m,k,t}-d^I_{m,k,t})}, \qquad (14)$$

where

$$\pi(f_{m,t}) = \Phi\left(\frac{c_m - \omega_m \cdot f_{m,t}}{\sqrt{1-\omega_m^2}}\right) \qquad (15)$$

is the homogenous conditional default probability of an infecting borrower in rating grade m. Due to the conditional independence of borrowers, the probability of observing $d^I_{k,t}$ defaulting firms within industry segment k is the product of the default probabilities in the m rating grades

$$P^I_{k,t}(\boldsymbol{f}_t) = P(d^I_{k,t}|\boldsymbol{f}_t) = \prod_{m \in I_t[k]} P^I_{m,k,t}(f_{m,t}), \tag{16}$$

where $\boldsymbol{f}_t = (f_{1,t}, \ldots, f_{M,t})$ denotes the vector of systematic risk drivers.

Let $C_t\,[m,k] = \{i|S(i) = m, K(i) = k, \text{contaminated}\}$ be the set of contaminated firms in industry segment k and rating grade m, where $C_{m,k,t} = N_{m,k,t} - I_{m,k,t}$ denotes their number and $N_{m,k,t}$ the total number of firms. Conditional on the common risk factor *and* the default frequency of the infectors we obtain the probability of observing $d^C_{m,k,t}$ defaulting contaminated firms as

$$\begin{aligned} P^C_{m,k,t}(\boldsymbol{f}_t) &= P\left(d^C_{m,k,t}|f_{m,t}, d^I_{k,t}(\boldsymbol{f}_t)\right) = \\ &= \binom{C_{m,k,t}}{d^C_{m,k,t}} \cdot \pi(f_{m,t}, d^I_{k,t}(\boldsymbol{f}_t))^{d^C_{m,k,t}} \cdot \left(1 - \pi(f_{m,t}, d^I_{k,t}(\boldsymbol{f}_t))\right)^{(I_{m,k,t} - d^C_{m,k,t})}, \end{aligned} \tag{17}$$

where

$$\pi(f_{m,t}, d^I_{k,t}(\boldsymbol{f}_t)) = \Phi\left(\frac{c_m - \omega_m \cdot f_{m,t} - \beta \cdot \frac{d^I_{k,t}(\boldsymbol{f}_t)}{I_{m,k,t}}}{\sqrt{1-\omega_m^2}}\right) \tag{18}$$

is the conditional default probability of the infected firms given the systematic factor and the default rate of the infectors.

Due to the rule of conditional probability where the joint probability of two events A and B is given by $P(A \cap B) = P(A|B) \cdot P(B)$, the joint probability of observing $d^I_{k,t}$ defaulting infectors and $d^C_{k,t}$ defaulting infected firms is

$$P^{IC}_{m,k,t}(\boldsymbol{f}_t) = P(d^I_{m,k,t}, d^C_{m,k,t}|\boldsymbol{f}_t) = P^I_{m,k,t}(f_{m,t}) \cdot P^C_{m,k,t}(\boldsymbol{f}_t). \tag{19}$$

Due to the conditional independence of the segments and rating grades we obtain the likelihood for the whole portfolio conditional on the random factors as

$$P^{IC}_t(\boldsymbol{f}_t) = \prod_{m,k \in I_t[m,k]} P^{IC}_{m,k,t}(\boldsymbol{f}_t). \tag{20}$$

Then the unconditional joint distribution is obtained by integrating over the joint distribution $g(f_{1,t}, \ldots, f_{M,t})$ of the systematic random factors

$$P^{IC}_t = \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \prod_{m,k \in I_t[m,k]} P^{IC}_{m,k,t}(\boldsymbol{f}_t)\, g(f_{1,t}, \ldots, f_{M,t})\, df_{1,t} \cdots df_{M,t}. \tag{21}$$

To the extent that all systematic factors are independent, this joint distribution drops to the product of the unconditional marginal distributions.

Finally, if we observe these default patterns for a whole time series of independent years the log-likelihood function is

$$l = \sum_{t=1}^{T} \ln \left\{ \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} \prod_{m,k \in I_t[m,k]} P_{m,k,t}^{IC}(\boldsymbol{f}_t)\; g(f_{1,t}, \ldots, f_{M,t})\; df_{1,t} \cdots df_{M,t} \right\}. \tag{22}$$

For a given default time series this function is optimized with respect to the parameters $c_1, \ldots, c_M$, $\omega_1, \ldots, \omega_M$, $\sigma_{2,1}, \ldots, \sigma_{M,M-1}$, and β.

5 Empirical Data

For the empirical analysis we use yearly bond default data for the years 1983 to 2003 published by Moody's Default Risk Service. In 1983, Moody's changed its rating methods resulting in a break in the time-series. Therefore, no data before 1983 is used. The dataset contains information regarding the respective face amounts of the exposures, the firms' countries of legal incorporation, the rating grades assigned by Moody's, and information on whether or not the bonds defaulted. Only bonds issued by US-American firms are used for our analysis. Altogether, we obtain a database with 30,643 firm-years including 820 defaults over the above-mentioned time horizon.

All firms are grouped into two sub-portfolios named "investment grade" (rating grades Aaa to Baa) and "speculative" (rating grades Ba to C), respectively. The total number of bonds varies over time due to matured, defaulted, and newly issued bonds. Figure 2 shows the number of bonds in the database per year.

The average default rate equals 2.68 percent, 0.16 percent, and 5.47 percent in the total portfolio, the investment grade portfolio, and the speculative portfolio, respectively. The time-dependency of the default rates can be seen in Figure 3.

All firms are assigned to one of 206 Moody's specific industries. The size of the industries in the database differs considerably, with the largest (U.S. bank holding) containing 1499 firm-years and the smallest (ship construction/repair) containing just one firm-year.

6 Results

6.1 Simulation Study

Before we present the empirical results of our estimation model we will demonstrate the robustness of the estimation procedure in a simulation study. Consider a synthetic portfolio consisting of 800 firms. 100 firms belong to the industry sector X, 200 to the industry sector Y, and 500 to the industry sector Z. It is assumed that within each industry sector

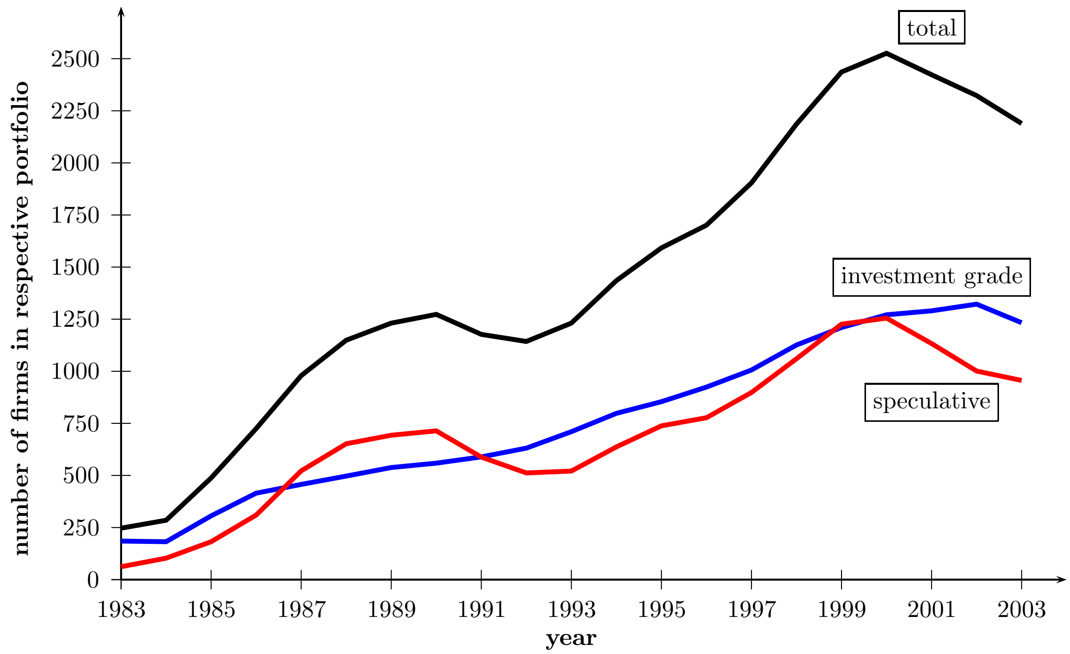

Figure 2: Portfolio sizes per year (empirical data)

20% of the firms may infect the remaining 80%, no matter whether or not the firms belong to the same rating grade. Within each industry-sector/contagion-group combination 50% of the firms belong to segment A (rating grade A) and 50% to segment B (rating grade B). Each segment is assumed to be homogeneous in terms of default probabilities and correlations. Figure 4 illustrates the number of firms in the respective industry-sector/contagion-group combinations for industry sector X.

Since the example takes just two segments into account, the covariance matrix of the systematic risk drivers reduces to

$$\mathrm{Cov}\left(f_A, f_B\right) = \begin{pmatrix} 1 & \sigma \\ \sigma & 1 \end{pmatrix}. \tag{23}$$

The assumed probabilities of default, the asset correlations, the contagion factor β as well as the covariance σ can be taken from table 1. The parameter values are chosen arbitrarily, different values lead to similar results.

	rating grade A	rating grade B
default threshold c	-1.64485	-1.28155
corresponding autonomous default probability π	0.05	0.1
asset correlation ρ	0.2	0.1
contagion factor β	-2	
covariance σ	0.5	

Table 1: True parameters in synthetic portfolio

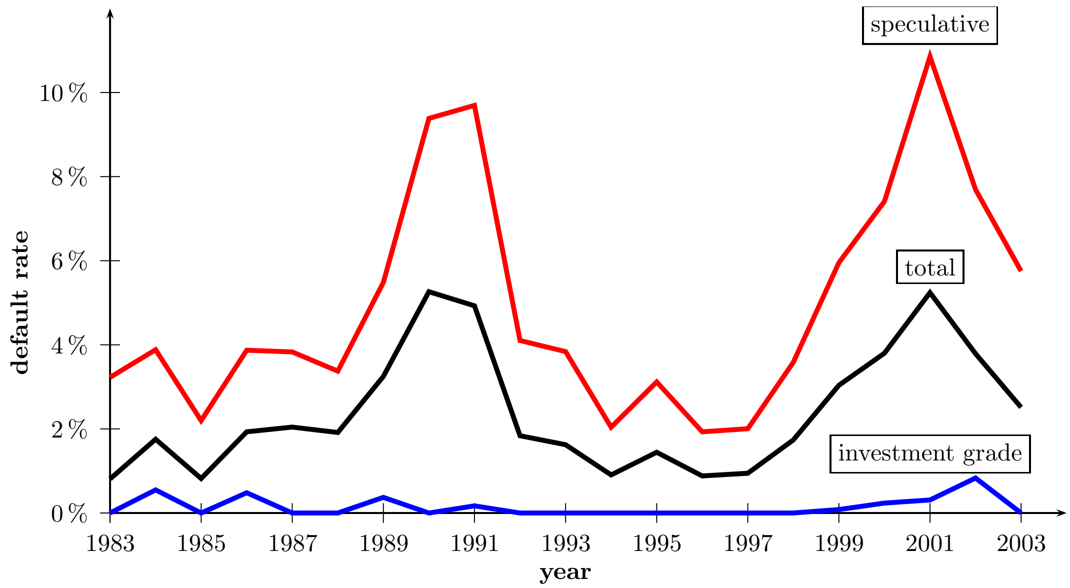

Figure 3: Yearly default rates (empirical data)

We now undertake a simulation study consisting of the following steps:

1. Generate a time-series of simulated defaults in the portfolio according to the stochastic model described by equations (8) and (12). The length of the simulated time series equals 20 years, which is also the length of the time-series used for the empirical analysis.

2. Estimate the parameters π^A, π^B, ρ^A, ρ^B, β and σ from the simulated default data using equation 22.

3. Repeat steps 1 to 2 10,000 times.

Table 2 summarizes characteristics of the estimated parameters.

	N	mean	std. dev.
autonomous default probability π^A	10,000	0.0499896	0.0121284
autonomous default probability π^B	10,000	0.1000272	0.0137923
asset correlation ρ^A	10,000	0.1891247	0.0543912
asset correlation ρ^B	10,000	0.0941505	0.0301767
contagion factor β	10,000	-2.0220358	0.2955764
covariance σ	10,000	0.4958245	0.1988554

Table 2: Parameter estimates (simulation study)

While the (autonomous) probabilities of default are estimated accurately, the asset correlations show a small downward bias. These results are in line with the findings of Gordy/Heitfield (2000). The estimation of the contagion factor β is of particular interest. As can be taken from table 2, the average estimated contagion factor ($\hat{\beta}$ = -2.022) is very close to the true value of -2. The major concern in modeling credit risk is the underestimation

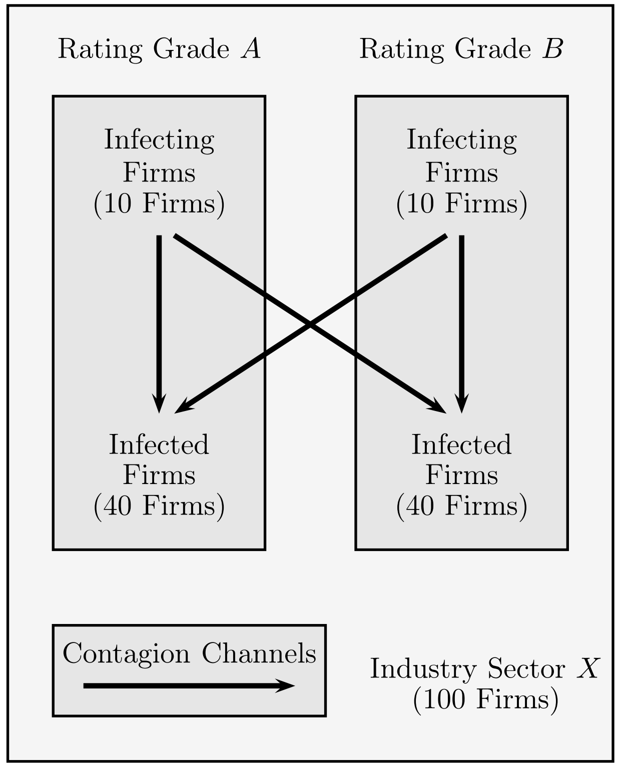

Figure 4: Contagion channels

of risk. The slight overestimation (in absolute terms) of the contagion factor, however, enhances the conservativeness of the resulting loss distributions and risk measures. The average estimated covariance ($\hat{\sigma} = 0.496$) falls just a little bit below the correct value of 0.5.

Analogous results are obtained using different parameter combinations for the default probabilities, asset correlations, the contagion factor, and the covariance. The PDs are always estimated accurately, the asset correlations show a small downward bias, the estimation effect (in absolute terms) is slightly overestimated, and the covariance exhibits a small downward deviation. All in all, the parameter estimation seems to work very well.

6.2 Empirical Estimation Results

We now turn to the analysis of the empirical data presented in chapter 5. First of all, some assumptions regarding the contagion channels have to be made. As already pointed out, we assume that credit contagion occurs within industry sectors. All firms in the database are assigned to one of 206 Moody's specific industries. We now assume that within each industry sector consisting of more than four firms the 20% biggest firms (measured by exposure size) may infect the remaining 80%, no matter whether or not the firms belong to the same rating grade. The intuition behind this assumption is that the default of a large bond may set signals for the entire industry sector. The ratio 20/80 is chosen rather arbitrarily and follows Pareto's principle which states that in many areas 20 percent of something are responsible for 80 percent of the results.

The resulting parameter estimates for the two segments "investment grade" and "speculative" can be taken from table 3. All parameter estimates are significant at the 1% level with the exception of the default probability and the asset correlation in the portfolio "investment grade". This is due to the low number of defaults in this subportfolio and does not challenge our model. The most interesting parameter is the contagion factor β, which equals -1.0020 and is highly significant as well. Obviously, contagion effects indeed play a role in credit portfolios and therefore should not be ignored when assessing credit risk.

	estimate	std. err.	p-value
default threshold (investment grade) c^{inv}	-2.9960	0.1553	<0.0001
default threshold (speculative) c^{spec}	-1.6629	0.05342	<0.0001
autonomous default probability (investment grade) π^{inv}	0.001368	0.000697	0.0644
autonomous default probability (speculative) π^{spec}	0.04817	0.005347	<0.0001
asset correlation (investment grade) ρ^{inv}	0.1872	0.1097	0.1042
asset correlation (speculative) ρ^{spec}	0.04747	0.01632	0.0090
contagion factor β	-1.0020	0.1643	<0.0001
covariance σ	0.7894	0.1733	0.0002

Table 3: Parameter estimates (empirical data)

6.3 Stress Testing

We now stress-test our portfolio with respect to the contagion factor β. First of all, using the maximum-likelihood estimates shown in table 3 we generate the loss distribution of the empirical portfolio for the year 2003. All exposures were set to 1 in order to focus on the effect of the varying contagion parameter on the loss distribution only. Ceteris paribus, we then generate loss distributions using different parameter values for the contagion factor in the range from β = -2 to β = -5. Figure 5 visualizes three of the resulting loss distributions. The dark blue bar represents the distribution using the ML-estimate of β = -1.0020, the light blue bar using β = -3, and the red bar using β = -5. By shifting the contagion factor the mean, the variance, and the weight in the tails are increased.

Table 4 contains the average default rates as well as the popular risk measures Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) for the different scenarios. The risk measures, visualized in figure 6, are defined as follows:

Definition 1 (Value-at-Risk) $\text{VaR}_\alpha(X) = \inf\{x \in \Re | P(X \leq x) \geq \alpha\}$.

Definition 2 (Conditional Value-at-Risk) $\text{CVaR}_\alpha(X) = \text{E}\{X | X \geq \text{VaR}_\alpha(X)\}$.

One can see the dramatic impact of a shift of the contagion factor on the high quantiles of the loss distribution. It is noteworthy that this shift affects the infected firms only. The loss distribution of the infecting firms remains completely unaffected. In contrast, the effect on a portfolio consisting of infected firms only is even greater.

Instead of stress-testing the portfolio with respect to the contagion factor β one could also stress the default rate of the infecting firms. Since the contagion factor β and the default rate

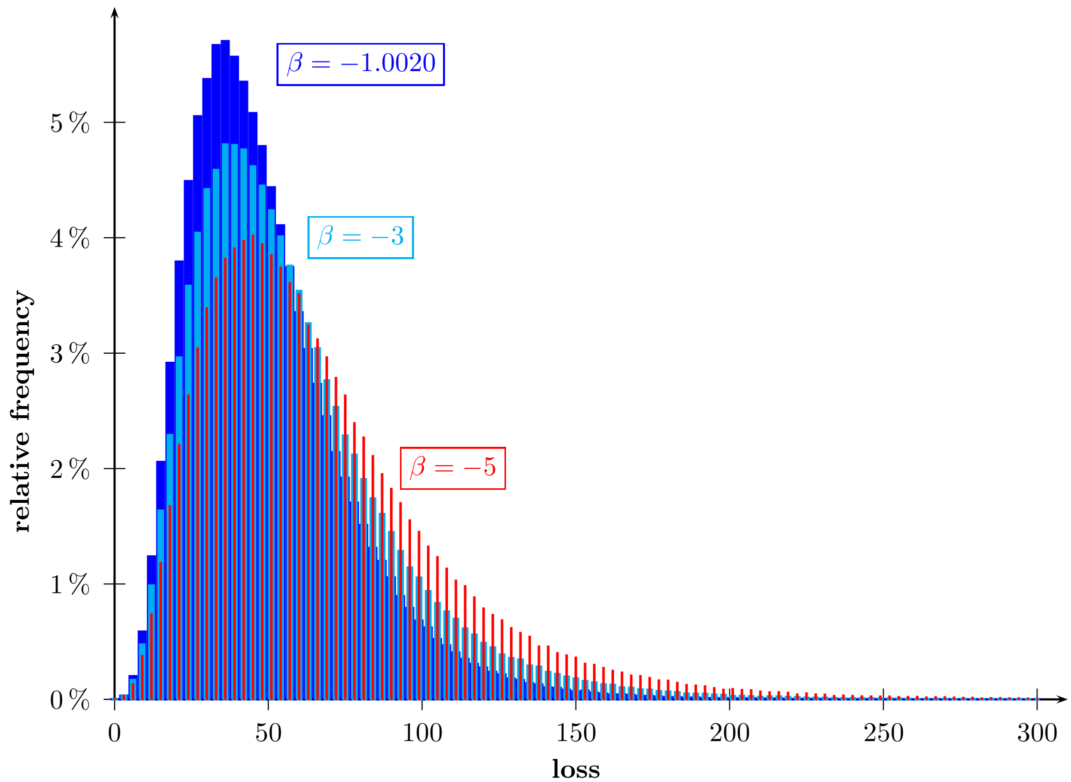

Figure 5: Loss distributions (stress testing)

	β = -1.0020	β = -2	β = -3	β = -4	β = -5
average default rate	2.28%	2.44%	2.63%	2.85%	3.09%
99%-VaR	139	154	170	189	207
99.9%-VaR	203	228	253	284	313
99.9%-CVaR	234.66	266.08	298.72	333.97	376.14

Table 4: Average default rate and selected risk measures of loss distributions

of the infecting firms appear in the model equations in a multiplicative way (see formula 12) the effect of stressing either factor is identical. Stressing the default rate of the infecting firms is similar to the approach by Neu/Kühn (2004) who propose to stress-test their model by setting specific firms into the default state.

7 Conclusion

In this paper, we have extended the standard factor model to allow for credit contagion. Contagion is assumed to occur within industry sectors. Every industry sector consists of infecting and contaminated firms. We assume a one-way dependency structure with the probabilities of default of the contaminated firms depending on the default rate of the infecting firms in the respective industry sector.

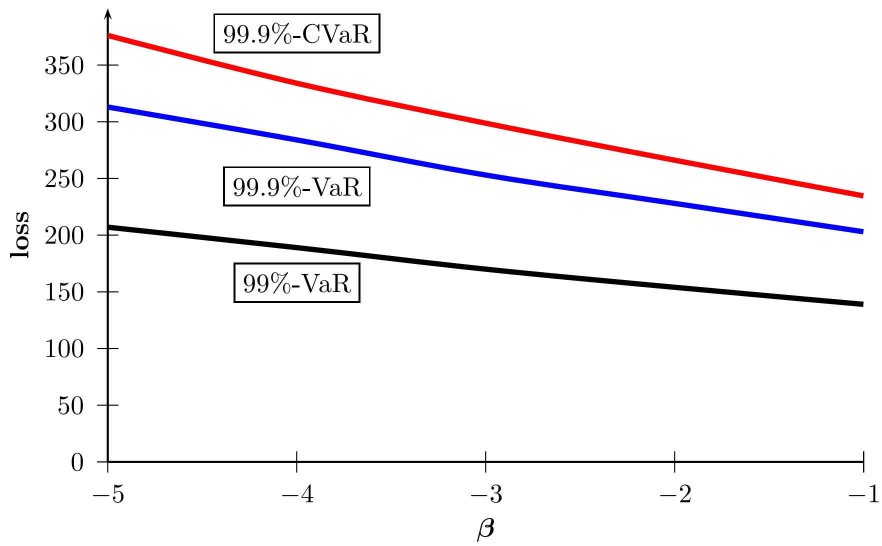

Figure 6: Selected risk measures of loss distributions

Our key findings can be summarized as follows: We have shown in a simulation study that contagion factors can be estimated accurately from historical data using a maximum-likelihood approach. Parameter estimation from empirical default data has revealed the existence of significant contagion effects in real-world portfolios. Concluding, we have stress-tested the empirical portfolio with respect to the contagion factor, demonstrating the dramatic impact of a shift of the contagion factor on various risk measures. Our findings are not only relevant for financial institutions but also for regulatory purposes. In order not to underestimate credit risk, contagion effects have to be implemented in credit portfolio models.

A possible extension of our model is the inclusion of segment *and* industry sector specific systematic factors. At the moment, the historical time-series of empirical default data are too short to estimate the necessary additional parameters. However, if more data becomes available, the estimation can be easily implemented.

References

Basel Committee on Banking Supervision 2004. International Convergence of Capital Measurement and Capital Standards – A Revised Framework, Bank for International Settlements, June 2004.

Bluhm, C., Overbeck, L., Wagner, C., 2003. An Introduction to Credit Risk Modeling, London.

Credit Suisse Financial Products, 1997. CreditRisk+ – A Credit Risk Management Framework, London.

Davis, M., Lo, V., 2001. Infectious defaults, Quantitative Finance 1, 382-387.

Egloff, D., Leippold, M., Vanini, P., 2004. A simple model of credit contagion, Working Paper, University of Zurich.

Finger, C., 1998. Sticks and Stones, RiskMetrics Group Working Paper, October.

Frey, R., McNeil, A., Nyfeler, M., 2001. Copulas and Credit Models, Risk 14, October, 111-114.

Giesecke, K., Weber, S., 2002. Credit contagion and aggregate losses, Working Paper, Cornell University and Technische Universität Berlin.

Gordy, M.B., 2000. A Comparative Anatomy of Credit Risk Models, Journal of Banking and Finance 24, 119-149.

Gordy, M.B., Heitfield, E.A. 2000. Estimating Factor Loadings When Ratings Performance Data Are Scarce, Board of Governors of the Federal Reserve System, Division of Research and Statistics, Working Paper.

Hamerle, A., Rösch, D., 2007. Parameterizing Credit Risk Models, forthcoming: Journal of Credit Risk 2(4).

Hamerle, A., Rösch, D., 2005. Misspecified Copulas in Credit Risk Models: How Good Is Gaussian?, Journal of Risk 8(1), 41-58.

Jarrow, R.A., Yu, F., 2001. Counterparty Risk and the Pricing of Defaultable Securities, Journal of Finance 56, 1765-1799.

Merton, R., 1974. On the pricing of corporate debt: the risk structure of interest rates, Journal of Finance 29, 449-470.

Neu, P., Kühn, R., 2004. Credit risk enhancement in a network of interdependent firms, Physica A 342, 639-655.

Pinheiro, J.C., Bates, D.M., 1995. Approximations to the Log-likelihood Function in the Nonlinear Mixed-effects Model, Journal of Computational and Graphical Statistics 4, 12-35.